Summary Findings of the Workshop on the Frontiers of Extreme Computing

Thomas Sterling (Louisiana State University)
Erik DeBendictis (Sandia National Laboratory)
David H. Bailey (Lawrence Berkeley National Laboratory)
9 November 2005

1. Motivation

Over the last few decades, the field of high-performance scientific computing has progressed on a Moore's Law exponential growth curve, driven by advances in semiconductor technology, scalable computer architecture, parallel numeric algorithms, and enabling system software. This development, which is without peer in the history of human technology, has catalyzed explosive growth not only in scientific computing, but also in the much broader world of information technology industry and e-commerce, with enormous economic and societal benefit to the U.S.

However, a detailed examination of future trends exposes some potential barriers to the continued expansion of this capability. These include:

- Increasing electric power consumption and heat dissipation challenges.
- Increasing difficulties in raising clock rates, due to various physical limitations.
- Increasing barriers to continued feature size reductions, due to basic barriers in photolithography and longer-term limits at the nanometer scale.
- Degradations in computational efficiency, due to the "memory wall."
- Exhaustion of architectural ideas for uniprocessor performance improvements.
- Decreasing hardware reliability, due to nano-scale feature sizes on one hand and increasing system sizes on the other.
- Decreasing software reliability, due to complex interactions in very large systems.
- Difficulties in programming very large scientific and commercial applications.
- Difficulties in understanding performance effects in large-scale systems.
- Looming barriers due to concurrency limits in applications (i.e., Amdahl's law).

On the other hand, if anything there is at present an explosion of interest in very large scientific applications, particularly in the arenas of national security, environmental protection, energy technology, nanotechnology and biotechnology. Important applications have been outlined that will require thousands or millions of times as much computational power and memory as is available on today's largest systems. In particular, there will be important requirements for systems beyond 1 Pflop/s (i.e., 10^{15} floating-point operations per second), which will be achieved by the end of the decade, extending to 1 Eflop/s (i.e. 10^{18} flop/s, also written 1 Exaflops) and even 1 Zflop/s (i.e., 10^{21} flop/s, also written 1 Zettaflops) or beyond.

To understand the nature of such potential systems and the potential scientific applications that might run on them, the Workshop on the Frontiers of Extreme Computing was conducted bringing together experts in these different domains to consider the medium and far future. This report briefly summarizes the findings of this workshop.

2. Objectives

The overall goal of the Workshop on the Frontiers of Extreme Computing was to establish the challenges, opportunities and roadmap for future directions in high-performance computing, in order to determine the necessary initiatives required to be undertaken to ensure continued capability growth in computing technology in the future. The primary objectives of this workshop are:

1. Establish application drivers and quantify the scale of requirements.
2. Project the International Technology Roadmap for Semiconductors (ITRS) to the end of Moore's Law (i.e., the end of conventional semiconductor technology).
3. Define barriers to continued capability growth in medium and long term.
4. Identify emerging concepts in device technology, parallel computer architecture, and system software and methods as potential approaches to surmount challenges.

3. Highlights

The workshop comprised four days of intense plenary and breakout working meetings, including prepared presentations and dynamic discussions to explore the range of topics and issues demanded by the workshop objectives, and to derive a set of consensus findings and recommendations.

The workshop identified a number of exemplary applications whose demands could exceed current computing capabilities by a factor of a million or more. Among these were the strategic problems of climate modeling and controlled fusion, both of which project future computing requirements in the range of 1 Eflop/s to 1 Zflop/s. Other problems that were data intensive rather than numeric intensive were also recognized, including a certain class of graph problems that are important for national security applications.

Given the enormous momentum of the global economy in information technology, ranging from semiconductor technology companies to device manufacturers to system vendors to software companies, those at the conference concluded that most likely the semiconductor technology companies will manage to continue their current rate of Moore's Law progress (with memory densities doubling every 18 months and CMOS clock rates increasing at a rate of 17% per year) for at least another five years and possibly ten years. This is also the conclusion of the International Technology Roadmap for Semiconductors (ITRS), a periodically updated assessment of the future prospects in the semiconductor industry. However, several significant challenges are looming. Most severe among these problems is the issue of power consumption, due to effects of dynamic transistor switching and static leakage currents. Other challenges are continuing the current pace of reduc-

tions in lithography feature size (which currently is near the limits of conventional lens-based systems), yield and reliability, and cost of fabrication.

Several future hardware technologies are emerging that may address one or more of these problems, while bringing other opportunities as well. Some of the more promising technologies include nanotube memories and transistors, nano-imprint crossbar devices, superconducting single flux quantum (SFQ) devices, and quantum dots. Taken collectively, these technology ideas offer the potential of very high speed, very fine grain scale, low power consumption, or some mix of these virtues in combination. For example, nanotube transistors are among the smallest switching devices devised, and SFQ devices have demonstrated the fastest switching speeds of any logic to date while operating at very low power. Some of these research ideas are fairly well along to practical deployment. For instance, Hewlett-Packard has already demonstrated nano-imprint crossbar memory devices with 17 nanometer feature sizes (five times smaller than the current semiconductor state-of-the-art, or 25 times smaller in area), and is confident that they can reduce this to four nanometers (20 times smaller in linear dimension, or 400 times in area) within two years. What's more, they have demonstrated that these devices can be used for logic operation as well as memory.

Given these several "back-up" technologies waiting in the wings (if not directly challenging the semiconductor industry), the consensus of the experts at the meeting is that almost certainly Moore's Law will continue at least through 2015, and possibly to 2020 and beyond, although the industry may have to undergo one or more somewhat disruptive paradigm shifts during this time period.

Looking further into the future, even these futuristic technologies exhibit identifiable limits of power and performance. Yet at least some strategic scientific applications have already been sketched to require more capability than such advanced devices could deliver. Thus, two additional strategies were considered at the workshop: reversible logic and quantum computing. Reversible logic provides the means to realize computational devices that operate below the Landauer limit, a fundamental boundary in energy consumption per mutable bit operation, below which no conventional computing device can perform. Reversible logic circumvents this limitation by recycling energy, in effect. However, utilizing reversible logic would require substantial reorganization of computing to avoid the loss of any state. Additional studies are needed to determine if some hybrid of reversible and conventional computing might be a workable solution. Along this line, the application of reversible logic to single flux quantum (SFQ) logic looks promising, and the workshop included a specific suggestion that the 1 Zflop/s performance level could be reached with reversible logic applied to quantum dots.

Quantum computing holds the most exciting potential of any technology considered at the workshop. This is an entirely new paradigm of computing which, for certain algorithms, can deliver polynomially or even exponentially faster execution times, compared with conventional computing, thus opening the way to solutions of certain problems that might never be approachable through other means. The list of potential applications for quantum computing presented at the conference was notably larger than once thought

possible. However, the consensus of those at the workshop is that even if quantum computing is ultimately realized and is able to perform this class of applications well, a large number of other applications is likely to remain that require classical approaches. In other words, quantum computing, even if ultimately successfully realized, will be a "niche" solution, and there will be a continuing need for classical computing.

While technology has been an important driver for the overall performance gain witnessed over the last two decades, other key factors have included application algorithms and computer architecture. In some cases, it has been estimated that algorithmic techniques have accounted for as much as half of the total performance gain observed. Computer architecture has been responsible for the exploitation of the device technology speed and density enhancements over the years while exploitation fine and coarse grain parallelism as well. However, the majority of conventional known architecture structures and mechanisms employed have been already applied, and there are very few ideas left to pursue, while retaining the basic von Neumann model to which essentially all such systems have adhered for the last 40+ years. New architecture strategies were proposed at the workshop including processor in memory, and cellular automata that could extend architecture performance yet further, but at the sacrifice of legacy application codes and architectures.

In the following sections, the main findings of the workshop are summarized along with some key recommendations to Federal funding agencies and the research community as a whole. In addition, some open issues for which a consensus view is not yet possible are provided.

4. Strategic Findings

4.1 Application Drivers

There are applications critical to the long term welfare of the nation and the world in science, technology, and national security that demand computing capability in the 1 Eflop/s performance regime and beyond for their potential to be fully realized. Such problems include climate modelling, controlled plasma fusion simulation, microbiology and drug synthesis, and world knowledge graph problems. System architecture imposes challenges to efficient large-scale application programming and execution including the memory wall (latency and bandwidth), the need to exploit ever more fine grain parallelism, and reliability.

4.2 Technology Directions and Disruptions

There is a continuum of computing approaches beginning with today's complementary metal-oxide silicate (CMOS) device technology and extending to progressively more powerful yet less developed ideas that match the needs of applications across the entire spectrum.

The semiconductor industry maintains a roadmap for CMOS, setting a target path for Moore's Law and identifying risks. CMOS is certain to maintain its current path for at least the next five years, and most likely for ten. As mentioned above, a multi-billion dollar market, driven by a worldwide information economy, has created a huge economic incentive for research and development in this area. Between five and 10 years out, there is significant risk that CMOS might not be to sustain the current rate of exponential performance gains. However, this risk is substantially mitigated by several classes of promising technologies, including nano-scale imprint devices, carbon-nanotube devices, magneto-resistive random-access memory (MRAM) and SFQ. Between 15 and 25 years in the future, a thermodynamic limit looms that may create a ceiling for the performance of logic as it is employed today. The only known long term solutions past this floor are reversible logic and quantum computing, as mentioned above.

It is clear that there is unlikely to be a single replacement for the semiconductor transistor but rather a set of complementing devices and methods that together will provide the necessary functionality at unprecedented density, speed, and low power. It is anticipated that some of these new technologies may even contribute to end-user system in the mid-term outlook (5-15 years). Along this line, increasing capital investment costs of new fabrication facilities will reduce the number of U.S. manufacturing lines. However, fault tolerance techniques will improve effective yield and mitigate these cost factors to some degree.

4.3 Architecture and Computing Models

It is expected that for the next few years, commodity clusters and massively parallel processors (MPPs), comprised of commercial off-the-shelf (COTS) multi-core commercial processor chips and dynamic random-access memory devices (DRAMs), will serve a significant part of the high-end workload. Part of the reason for this rather conventional approach is that the computer architecture research pipeline is essentially empty, with little or no Federal funding for academic research during the past few years in innovative architectures. Quite a number of academic researchers and young PhD graduates who once were interested in the field have gone elsewhere, such as to the booming networking community. At the same time, conventional techniques that have sustained the past growth in architecture performance such as pipelining, ILP, branch prediction and speculative execution, have been exhausted. The new reliance on multi-core processor designs is imposing more challenges than solutions.

Fortunately, there are a few credible ideas for future architectures that appear promising, both to deal with vastly greater parallelism as well as overcoming performance degradation factors such as the memory wall, overhead, and starvation. However, these innovative ideas need to be further developed and brought to the stage of real prototype demonstrations. In addition, both active power management and fault tolerance must become mainstream architectural features in the 1 Eflop/s domain and beyond. Future architectures may rely on increased specialization and heterogeneity of structure to achieve performance goals for higher efficiency of key algorithms.

4.4 Programming Models, Languages and Tools

Programming languages and environments must also evolve if the key challenges of memory wall, concurrency, and reliability are to be addressed. But it is possible that such innovation can be achieved through augmentation of existing programming models and frameworks rather than demanding entirely new disruptive languages and techniques. Scientists developing large-scale computations over the last decade using common tools (such as Fortran-90, C++ and MPI) have produced a critical mass of portable scalable application packages. These packages, which are based on fairly standard geometric decompositions used in many applications, provide a route for the evolution of codes in terms of algorithms, software, and languages. Nonetheless, the available user programming models and related software facilities are fragile both in terms of reliability and performance. This will require significant improvements to enhance system efficiency and programmer productivity. Use of unstructured methods (for example with adaptive mesh refinement) will require more advanced methods for data representation, distribution, and decomposition, as well as finer control of data parallelism. While independent software vendors now support distributed memory programming paradigms, much more is required of the software industry to ensure wide adoption and quality implementations of parallel applications.

5. Open Issues

While reversible logic may provide the only classical path beyond Landauer's limit, such a technology change would be quite disruptive. In particular, reversible logic does not appear to be compatible with conventional CMOS and would require alternative formulations of basic logic. While it is notable that both IBM and NSF agreed at this workshop to support additional research into reversible logic, the path forward has much uncertainty.

Another unresolved issue is how to fully engage independent software vendors to support high-end applications and system software. Independent software vendors are not thinking at anywhere near the huge scales or across a wide range of applications and algorithms. This problem is aggravated by several issues including:

- Few participants from academia have access to the state-of-the-art high-end systems.
- The validation and verification problem for systems and scientific applications is increasingly unavoidable, yet little has been done to address it.
- Legacy applications and newly written applications may need different solutions.

Along this line, the IEEE 64-bit floating-point arithmetic standard used in present-day computing is widely expected to be inadequate for these future systems and their demanding applications; indeed, problems in this arena are already surfacing. Thus vendors need to seriously consider incorporating hardware support for the IEEE 128-bit standard, and both application scientists and software vendors need to prepare to support this datatype.

New architectures will require custom software to take advantage of their specialty properties and features. For heterogeneous systems incorporating specialty technology or elements of a discontinuity nature, special software will have to be developed such as for quantum coprocessors, field-programmable gate arrays (FPGAs), processor-in-memory units (PIMs), cellular automata, and others.

As mentioned in the introduction, one looming obstacle that may prove fatal for some types of high-end applications is the worry that these applications will not be able to exhibit (or that systems will not be able to exploit) sufficient levels of concurrency to keep multi-hundred-thousand-CPU systems busy, thus suffering from severe Amdah's Law reductions in performance. Such challenges may have to be addressed by requiring codesign methodologies between hardware and supporting software. It is an open question as to what are the minimal extensions to MPI, C, and Fortran to exploit unique features not unlike those for vector or multithread extensions to Fortran.

There was some debate as to whether a truly general-purpose single-class architecture could be developed. While there are many cogent arguments supporting the heterogeneous ensemble of specialty processors strategy, a minority view expressed is that a general-purpose system is possible assuming the hardware elements are designed expressly for the purpose of operating in synergy with other elements for optimal concurrency. This has not been tried in many years.

6. Recommendations

It was universally agreed that the nation needs to re-establish sustained funding support for more than one community of computer architectures, system software, and programming models, so that these communities can explore designs that are more than incremental improvements to existing product lines. This is consistent with National Research Council (NRC) Recommendation #6. It also appears in two recent Defense Science Board (DSB) studies.

Each of the following three hardware directions must be supported with significant new research sponsorship, in an effort to match both the technology time windows and the diversity of application workload requirements. These directions include:

- Commodity architectures following Moore's Law to at least 2020.
- Specialty and custom architectures.
- Discontinuities (e.g., quantum computing).

It should be emphasized that several areas are included in the national HPC systems research portfolio that are not currently deemed important, including novel computer architecture, new programming models, reversible and superconducting logic, and a plethora of nano-scale technologies. Also, there needs to be continued effort at defining applications problems suitable for quantum computers, beyond the realm of cryptographic problems that are presently the primary focus of funding.

Federal agencies responsible for supercomputing should underwrite a community effort to develop and maintain a roadmap for both hardware and software. This is consistent with NRC Study Recommendation #5, and can be undertaken immediately through mutual agreement among program implementers.

Application roadmaps with quantitative cost and science milestones should be developed for key problems critical to the long term welfare and security of the nation. These roadmaps should be used to motivate and guide national policy of investment in high risk/high payoff research for future HPC systems development.

The academic research community must adapt a more progressive attitude concerning areas of research meriting recognition and coverage in community led forums, journals, and proposals under peer review to encourage young researchers to rejoin the high-end computing field.

7. Conclusions

In summary, the field of high performance computing has entered a period of significant uncertainty, with an end to Moore's Law for conventional semiconductor technology, and the confidence of continued exponential growth of performance through conventional approaches dimming beyond five years into the future. However, promising new technologies are emerging that could sustain the growth (given sufficient and timely down to nanometer scale for another decade or longer. Ultimately, atomic scale and fundamental power limits will demand a radical departure from conventional logic practices potentially requiring novel techniques based on reversible logic and quantum computing to be adopted. In all stages of this evolution, a renewed and aggressive investment in risky research is required of the Federal government if U.S. industry is to continue to lead the world in opportunity and capability.